# Shell Occupancy in RealQM: the Two Constructions Bracket the Exact Energy

Is the blending a state or a parameter of the electron model?

Claes Johnson[1]

[1]Royal Institute of Technology, Stockholm (emeritus)

**Abstract**

RealQM represents $N$ electrons by $N$ unit charge densities on non-overlapping domains of ordinary space, with the energy a localisation term plus the Coulomb energy of the configuration, pairs counted once. For a shell holding more than one electron the construction admits two readings, and both are in use: the electrons may be *split* into disjoint angular sectors, or *homogenised* into one shared density with the self-repulsion reduced by $(N_s-1)/N_s$. The first is what the stated non-overlap condition prescribes; the second is what the accurate atomic calculations use. On helium they give $-3.018$ and $-2.840$ Hartree against an exact $-2.9037$, so the choice is worth about 6% and nothing in the theory fixes it.

We connect the two constructions by a one-parameter family and locate the exact answer inside it. The connection rests on an identity: at an interface which is a plane of symmetry, the kinetic energy of the *sum* $\psi=\psi_1+\psi_2$ equals the sum of the two domain kinetic energies computed with the natural condition on each domain — to machine precision. Both ends of the family therefore carry the *same* kinetic energy, and the entire difference between them is the pair repulsion, which is a bounded functional. Parameterising the repulsion rather than the densities removes an obstruction: a family built on the densities sends the localisation cost to infinity at the split end, because a jump in $\sqrt{\rho}$ is not in $H^1$, and the split end is then unreachable.

With the kernel softening calibrated on the *one-electron* exact energy $-Z^2/2$, so that no two-electron information enters, the family reproduces both published endpoints of helium to 0.2–0.3%, and the exact two-electron energy lies strictly inside the bracket for $Z=2,3,4$ at both of two mesh spacings. Where along the family it lies is *not* resolved: the fraction drifts by 15–30% between the two meshes and the drift grows with $Z$, so no quantitative position is claimed here.

The energy is monotone along the family, and what that means turns on a question the formulation does not settle: whether a partially blended configuration is admissible as a *state*, or whether the blending is a *parameter* of the electron model. If a state, minimising selects full separation, the model over-binds by 4.2%, and the more accurate homogenised shell is a higher-energy configuration whose accuracy is cancellation rather than correctness. If a parameter, there is nothing to minimise over, and what the computation reports instead is that the required blending is *interior* — so RealQM would carry one *universal* constant, against a core radius per element or a functional per problem. The numbers are identical under both readings. The stated non-overlap condition appears to exclude blended states, but it also excludes the blended shell the accurate calculations use, so the formulation is not self-consistent here. We do not measure the constant: the position along the family is not converged, and whether one value serves different systems is untested.

## 1 The ambiguity

The RealQM energy for $N$ electrons on domains $\Omega_i$ is a localisation term plus the Coulomb energy with pairs counted once and no self-interaction. For one electron per domain the prescription is unambiguous. For a shell holding $N_s > 1$ electrons it is not, and two readings are in use.

The *split* reading follows the stated non-overlap condition: the $N_s$ electrons occupy disjoint angular sectors of the shell, two of them as hemispheres, four as a tetrahedron, eight as two interleaved tetrahedra. The *homogenised* reading gives all $N_s$ electrons one shared spherically symmetric density and reduces the self-repulsion by $(N_s-1)/N_s$, which is pair counting restated. The accurate atomic calculations use the second; the stated condition prescribes the first.

On helium the difference is not small. The two constructions give $-3.018$ and $-2.840$ Ha against an exact $-2.9037$, so the choice is worth about 6% of the correlation-sized part of the energy and straddles the answer. Nothing in the formulation selects between them.

## 2 The two ends share their kinetic energy

The obstruction to connecting them is apparent rather than real, and removing it is the technical content of this note.

Write $\psi = \psi_1 + \psi_2$ for the sum of the two one-electron functions. In the split construction each $\psi_i$ is supported on its own half and is discontinuous across the interface, so $\psi_i \notin H^1$ and its localisation energy is not defined as a distribution. But the sum is smooth, and the quantity actually computed in a split calculation — the localisation energy assembled over each domain with the natural condition on its own boundary — is exactly the localisation energy of that sum. We verified this numerically on an axisymmetric grid: for a state symmetric in $z$ with the interface on a cell face,

$$T[\psi] = 1.40322755, \qquad T_1^{\mathrm{dom}} + T_2^{\mathrm{dom}} = 1.40322755,$$

agreeing to $2.2 \times 10^{-16}$, and likewise at a second spacing. For a state deliberately skewed so that $\partial\psi/\partial z \neq 0$ at the interface the two differ by 1–2%, which is the straddling-face term that the symmetry kills. So at a plane of symmetry the split construction's kinetic energy is not an approximation to anything: it is the kinetic energy of one smooth function, attributed by region.

Two consequences. First, no state drops to zero anywhere and no $H^1$ condition is violated, because $\psi_1$ and $\psi_2$ are never formed separately. Second, and this is what makes the connection possible, both ends of the construction carry the *same* kinetic energy for a given total density:

$$\text{homogenised: } \rho_i = \psi^2 \;\Rightarrow\; T = \int |\nabla\psi|^2, \qquad \text{split: } \rho_i = 2\psi^2 \mathbf{1}_{\mathrm{half}} \;\Rightarrow\; T = \int |\nabla\psi|^2.$$

The entire difference between the constructions is the pair repulsion.

## 3 The family, and why it must be built on the repulsion

Let $w$ be a weight with $w(z) + w(-z) = 1$ and set $\rho_1 = 2w\psi^2$, $\rho_2 = 2(1-w)\psi^2$ with $\int \psi^2 = 1$. Taking $w = \frac{1}{2}$ gives the homogenised shell and $w$ a step gives the split shell. Holding the localisation and nuclear terms at their common value,

$$E(\lambda) = \int |\nabla\psi|^2 - 2Z\int \frac{\psi^2}{r} + U(\lambda), \qquad U(\lambda) = \int P[\rho_1]\,\rho_2, \qquad w = \left(1 + e^{-z/\lambda}\right)^{-1}, \quad (1)$$

minimised over $\psi$ at each $\lambda$. Both ends are finite: $U \to J[\psi^2]$ as $\lambda \to \infty$ and $U \to$ the cross-term between disjoint halves as $\lambda \to 0$.

It matters that the family is built on $U$ and not on the densities. A family built by partitioning $\sqrt{\rho}$ carries a localisation cost $\frac{1}{2}\int \psi^2 Q$ with $Q = \sum_i |\nabla\sqrt{2w_i}|^2$, which diverges as $w$ approaches a step — a jump in $\sqrt{\rho}$ is not in $H^1$. We computed that family as well; its energy rises monotonically away from full overlap, the split end is unreachable, and the exact answer lies outside it altogether. The apparent bracketing of the exact energy is recovered only when the parameterisation is placed on the bounded functional.

## 4 The exact energy lies inside the bracket

The kernel is softened, and the softening is calibrated *per element on the one-electron exact energy* $-Z^2/2$, so that no two-electron quantity enters the calibration and the two-electron result is a prediction. On helium the family then returns $-2.840$ at the overlap end against the published homogenised $-2.845$, and $-3.018$ at the split end against the published $-3.026$: both ends reproduced to 0.2–0.3% by a calibration that used neither.

| $Z$ | $h_0$ | overlap | exact | split | inside? | fraction along |
|---|---|---|---|---|---|---|
| 2 | 0.22 | −2.840 | −2.904 | −3.018 | yes | 0.358 |
| 3 | 0.22 | −7.198 | −7.280 | −7.478 | yes | 0.293 |
| 4 | 0.22 | −13.555 | −13.656 | −13.940 | yes | 0.262 |
| 2 | 0.16 | −2.852 | −2.904 | −3.022 | yes | 0.304 |
| 3 | 0.16 | −7.219 | −7.280 | −7.493 | yes | 0.221 |
| 4 | 0.16 | −13.586 | −13.656 | −13.964 | yes | 0.183 |

Two things are stable and one is not.

*The bracket holds.* At both spacings and for all three ions the exact energy lies strictly between the two constructions. And the endpoints improve under refinement: at $h_0 = 0.16$ helium's overlap end is $-2.852$ against the published $-2.845$ and its split end $-3.022$ against $-3.026$, agreeing to 0.25% and 0.13% respectively, by a calibration that used neither.

*Where along the family the exact answer lies is not resolved.* The fractions fall by 15%, 25% and 30% between the two spacings, the drift growing with $Z$, and their spread across the series *widens* under refinement rather than narrowing. A fraction is a difference of differences and is accordingly the most mesh-sensitive quantity here. An earlier draft of this note read the coarser row as a near-constant position of about one third and offered it as a candidate rule; the finer row does not support that and the claim is withdrawn. What the computation establishes is the bracket, not a position inside it.

The scaling is physically sensible. $\lambda$ is the length over which the two territories blend, so it is an avoidance length; the atom shrinks as $1/Z$ while $\lambda^*$ shrinks more slowly, so the blending region grows relative to the atom as $Z$ rises — which is what one expects when the nuclear attraction increasingly dominates the mutual repulsion and the electrons have less reason to avoid one another.

## 5 What the two ends are, and how wide the bracket is

The bracket is not a span of unexplained size. Both of its ends can be identified, and its width follows a law.

**The overlap end is the correlation energy.** For a closed two-electron shell the homogenised functional *is* Hartree, and Hartree for a closed two-electron shell is Hartree–Fock. So the overlap end should miss exactly the correlation energy, and it does:

| | overlap − exact | correlation energy | |
|---|---|---|---|
| He | +0.0426 | 0.0420 | (−2.903724 against HF −2.861680) |
| $H_2$, $R = 1.4$ | +0.0409 | 0.0408 | (−1.174476 against HF −1.133630) |

to four figures in both cases. The overlap end therefore has a known target and hits it, which is what makes the other end informative: the split end over-binds by 2.9 times the correlation energy for helium and 2.5 times for $H_2$, so the full bracket is between three and four correlation energies wide.

**The fractional width goes as $1/Z$.** By Section 2 the two ends carry the same kinetic energy, so the width is *entirely* pair repulsion. For a two-electron shell the orbital size goes as $1/Z$ and hence the repulsion as $Z$, while the total energy goes as $Z^2$. The width relative to the total should therefore fall as $1/Z$ — and it does, across every row computed, including $H_2$ as the $Z = 1$ member:

| | $Z$ | width | width/$\lvert E\rvert$ | 11.6%/$Z$ |
|---|---|---|---|---|
| $H_2$, $R = 1.4$ | 1 | 0.1414 | 12.04% | 11.60% |
| He | 2 | 0.1654 | 5.70% | 5.80% |
| $Li^+$ | 3 | 0.2800 / 0.2740 | 3.85% / 3.76% | 3.87% |
| $Be^{2+}$ | 4 | 0.3850 / 0.3780 | 2.82% / 2.77% | 2.90% |

The fit holds to a few tenths of a percentage point on six independent rows, at both mesh spacings, with one free constant. Two consequences are worth stating. The ambiguity is *worst at small $Z$*: hydrogen and helium are the extreme cases rather than representative ones, and by beryllium the choice is already worth under 3%. And the law is a prediction for the rest of the table, since it comes from the scaling of the repulsion rather than from any fitted quantity.

**What this does not extend to.** Every row above is a two-electron shell, where the split is a plane and Section 2's identity applies. A shell of four or eight is split by a tetrahedron or a cube, which has no single symmetry plane, so there is no reason for the kinetic energies to agree between the two readings there. If they do not, the width is no longer pure repulsion and the $1/Z$ derivation lapses with it. That is the first thing to test in extending this.

## 6 Is the blending a state or a parameter?

$E(\lambda)$ is monotone along the family: for helium, −2.845, −2.852, −2.868, −2.893, −2.929, −2.962, −2.990, −3.006 as $\lambda$ decreases. What that means depends on a question the formulation does not settle, and the two answers lead to quite different places.

**If the blending is a state.** Suppose configurations of different $\lambda$ are different states of one model, so that the ground state is obtained by minimising over $\lambda$ along with everything else. Then monotonicity decides it: the minimiser is the fully split configuration, and it over-binds helium by 4.2%. On this reading strict non-overlap is what the variational principle selects, the model over-binds, and the homogenised shell — the more accurate of the two and the one behind the published atomic totals — is not a variational competitor at all but a *higher*-energy configuration lying nearer measurement. Its accuracy would then be cancellation rather than correctness.

**If the blending is a parameter.** Suppose instead that the degree of blending is part of what an electron *is* in this theory, fixed once and not optimised per system. Then configurations of different $\lambda$ are different models rather than different states, there is nothing to minimise over, and one minimises the densities and interfaces at fixed $\lambda$. Nothing in the formulation forbids this, and it is the reading under which the practice of the accurate calculations — which use a blended shell throughout — is consistent rather than anomalous.

**Which reading the formulation supports.** The two differ precisely on whether a partially blended configuration is *admissible as a state*. If it is, a lower-energy admissible configuration exists and the ground state is there; if it is not, $\lambda$ is a constant of the theory. The stated non-overlap condition appears to exclude blended states, which would make $\lambda$ a parameter — but then the condition also excludes the blended shell that the atomic calculations use, so the formulation is not self-consistent on the point as it stands. Settling it is a matter for the formulation, not for computation.

**What is at stake.** The parameter reading is the stronger of the two, because the parameter would be *universal*: one number for all electrons, as against a core radius per element, or an exchange–correlation functional chosen per problem. On that reading the computation reported here says something positive rather than negative — that the required blending is *interior*, being neither full overlap nor full separation, so a single such constant is at least not excluded. On the state reading the same computation says the model over-binds and the accurate construction is fortuitous. The numbers are identical; the two readings disagree about what they mean.

What the computation does not do is measure the constant. That requires the position along the family, which §4 shows is not converged, and it requires the same value to serve different systems — which is the question the parameter reading turns on, and which is untested here.

## 7 What these numbers do not establish

Two spacings only. They suffice to show that the position along the family is not converged — which is why no position is claimed — but not to extrapolate one. Whether the fractions tend to a limit, and whether that limit is common across $Z$, is open and needs a third spacing at minimum.

The kernel is softened rather than bare. The softening is calibrated on a one-electron quantity, which is the cleanest available choice and makes the two-electron numbers predictions, but it is a regularisation and the absolute energies carry its error.

Only two-electron shells are treated. Whether the same position along the family describes a shell of four or eight — where the split is a tetrahedron rather than a plane, and where the

capacities 2 and 8 of the pair-and-octet construction are at stake — is untested here, and is the obvious continuation.